

# METHOD FOR A NETWORK DEVICE INSERTED BETWEEN POINT TO POINT CONNECTED STATIONS TO AUTOMATICALLY NEGOTIATE COMMUNICATION PARAMETERS BETWEEN THE STATIONS

5

## COPYRIGHT NOTICE

Contained herein is material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction of the patent disclosure by any person as it appears in the Patent and Trademark Office patent files or records, but  
10 otherwise reserves all rights to the copyright whatsoever.

## BACKGROUND OF THE INVENTION

### Field of the Invention

The present invention relates to data networking. In particular, the present  
15 invention relates to a method and apparatus for a network probe to synchronize enhanced modes of operation between two nodes connected via a point-to-point link, into which link the network probe is inserted.

### Description of the Related Art

20 With reference to Figure 1, prior art data networks generally utilized one or more shared media hubs, e.g., hub 10. Multiple nodes, or end user workstations, e.g., workstations 1, 2 and 3, were coupled to a shared communications that was, in turn, coupled to a port on the shared media hub. The hub 10 had multiple ports (e.g., ports 11, 21, 31), each coupled to a different shared communications medium. High end  
25 workstations, or servers, such as file servers or print servers, were also coupled via a dedicated or shared communications medium to a port on the shared media hub.



As the applications running on these data networks became more mission critical, and bandwidth utilization of the shared communications media increased, it became advantageous to monitor, for example, the performance and the error rates of data traffic on the shared communications media to facilitate proper operation of the data network. To accomplish this monitoring, network monitoring devices were configured into shared media hubs, or coupled to the port (41) of a shared media hub via a communications medium (40) as stand-alone devices (e.g., probe 4). In either configuration, the monitoring devices were typically referred to as probes. The probes would promiscuously monitor the data traffic on all shared communications media in the network and look at, for example, performance and error statistics, data traffic patterns and typical data flows across the shared communications media.

As shown in Figure 2, as performance requirements of prior art data networks continued to increase, and additional performance intensive applications were employed, the shared communications media coupled to the shared media hubs were typically divided into multiple network segments (e.g., network segments 201, 202 and 203) to reduce data traffic on each segment, although all network segments were still in the same collision domain, i.e., the network segments were not electrically isolated. Data communication between these segments generally utilized well known backbone, rather than switching, technology.

As performance requirements continued to increase to meet traffic demands, switches such as switch 220 illustrated in Figure 2 were used to segment the network into multiple collision domains. Segmenting the network into multiple collision domains prevented a data packet from one segment (e.g., segment 201) traversing the network to another segment (e.g., segment 202) unless the data packet was destined to a particular



device on the other segment. Such a determination was based, for example, on a destination address specified in the data packet.

The problem, however, in monitoring network performance in such an environment utilizing probes was that a single probe was required for each segment in order to promiscuously monitor the data traffic on that segment. With reference to Figure 3, as the data networks became highly segmented, it became evident that it was impractical to attach a probe to each segment in the network to promiscuously monitor all traffic. Rather, network administrators tended to concentrate probing activities to highly concentrated server farms or segments in the network where the traffic was the busiest, for example, a segment from a switch to a file server. These file servers were typically coupled via a dedicated point to point communications medium to a port on a switch to provide, for example, a data communications rate of 10 megabits per second, 100 megabits per second, or even 1000 megabits per second, to the file server. Connecting the file server using a dedicated point to point communications medium to the switch formed a single station network segment. In a single station network segment, it was impossible to attach a probe to that segment to promiscuously monitor network traffic because only a single port was necessarily available for coupling the segment to the switch. To overcome this limitation, a multiport repeater was inserted between the switch and the file server, e.g., repeater 233 between workstation 3 and switch 220 in Figure 3, thereby providing additional ports (on the inserted multiport repeater) to facilitate connection of a probe (e.g., probe 235) into the segment.

Although switch 220 in Figure 3 shows only six ports for purposes of illustration, it is understood that a switch may have sufficient ports to support, for example, ten, twenty, or more servers. In such a situation, it becomes impractical to attach a repeater between every server and switch port to promiscuously monitor data traffic, due to the



increased cost, space, and asset management responsibilities encountered as a result of the additional equipment. In addition, for each communications medium that was to be monitored, that network segment would have to be taken down, the server disconnected from the switch, the repeater inserted into the communications medium, and the server

- 5 communication reestablished. This process would be highly disrupting to data communications in the network. Moreover, in attempting to diagnose a performance problem, one would be required to shut down the network segment, insert a repeater, and couple a probe to the repeater in order to collect monitoring data. By the time the probe was operable, the performance problem may well have disappeared.

10

SECRET 200609160



With reference to Figure 4, a prior art probe as may be utilized in a typical switched data networking environment is illustrated. Multiple network segments 201, 202 and 203 are coupled to a shared media hub 200. Each segment is connected to separate modules 204, 205 and 206, respectively, within the hub. Each of the hub  
5 modules are coupled via a dedicated communications medium 401, 402 and 403 to an individual port 404, 405 and 406 on a switch 220. (Alternatively, segments 201, 202 and 203 may each be a dedicated communications medium, in which case, the segments would be directly coupled to respective ports 404, 405 and 406 on switch 220.) Ports on the switch are additionally shown connected either to a dedicated network device, e.g.,  
10 device 2 (perhaps an end user workstation or a server), or connected to prior art probe 400. More specifically, a port (e.g., port 407) on the switch may be connected to a port (e.g., port 408) on the probe. Another port (409) on probe 400, in turn, is coupled to a network device such as workstation 1. While the network device illustrated is a workstation, it is appreciated that the network device may be another switch, a server, or  
15 other network device. Probe 400 includes circuitry for repeating data packets between the switch and the network devices coupled to the probe.

The probe 400 utilizes internal bypass circuitry in promiscuously monitoring the communications medium coupling network devices 1 and 3 to the probe. The probe  
20 promiscuously monitors all traffic between the hub 200 and switch 220 destined for or received from either network device 1 or network device 3. All data traffic is captured, and potentially saved, e.g., for some form of analysis or statistical compilation. The probe analyzes those packets according to, for example, the remote monitoring standards RMON I or II. These standards promulgate, for example, specific statistical  
25 characteristics, such as user history, performance and error rates and traffic patterns between different workstations on the network across all layers of the International







[illegible]

5



## BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the following figures. Like references indicate similar elements, in which:

5

**Figure 1** is an illustration of a data network utilizing a shared communications media hub.

10

**Figure 2** is an illustration of a data network utilizing a switching hub (switch) to provide a dedicated communications medium to a network device.

**Figure 3** is an illustration of a prior art method of monitoring performance in a data network utilizing a repeater coupled to each network device.

15

**Figure 4** is an illustration of a prior art network probe.

**Figure 5** is a diagram of a finite state machine embodying the present invention.

20

**Figure 6** is a block diagram of a network probe in accordance with an embodiment of the present invention

**Figure 7** is a block diagram of a network probe configuration according to one state of the finite state machine depicted in Figure 5.

25

**Figure 8** is a block diagram of a network probe configuration according to one state of the finite state machine depicted in Figure 5.



[illegible]



## DETAILED DESCRIPTION OF THE INVENTION

The present invention provides a method and apparatus embodied in a network probe inserted between two nodes connected via a point-to-point link for negotiating and synchronizing enhanced modes of operation between the two nodes. In the following  
5 description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one of ordinary skill in the art that the present invention may be practiced without these specific details. In other instances, well-known architectures, circuits, and techniques have not been shown to avoid unnecessarily obscuring the present invention.

10 In alternative embodiments, the present invention may be applicable to implementations of the invention in integrated circuits or chip sets, wireless implementations, switching systems products and transmission systems products. For purposes of this application, the terms switching systems products shall be taken to mean  
15 private branch exchanges (PBXs), central office switching systems that interconnect subscribers, toll/tandem switching systems for interconnecting trunks between switching centers, and broadband core switches found at the center of a service provider's network that may be fed by broadband edge switches or access multiplexors, and associated signaling, and support systems and services. The term transmission systems products  
20 shall be taken to mean products used by service providers to provide interconnection between their subscribers and their networks such as loop systems, and which provide multiplexing, aggregation and transport between a service provider's switching systems across the wide area, and associated signaling and support systems and services.

25 With reference to Figure 5, a diagram of a finite state machine 500 embodied by the present invention is provided. The finite state machine 500 negotiates and synchronizes enhanced modes of operation between two nodes connected via a point-to-



point link. The finite state machine 500 is implemented at the network probe inserted between the two nodes to negotiate with each node between which it is inserted in accordance with the Auto-Negotiation function.

5 Pursuant to IEEE Std 802.3u, clause 28, the Auto-Negotiation function is used as the signaling method to automatically configure a preferred mode of operation between two CSMA/CD compatible network devices that use the 8 Pin Modular connector specified in ISO/IEC 8877 and that also encompass multiple operational modes. Annex IEEE Std 802.3u, draft 5.3, June 12, 1995, section 28B.3 Priority Resolution, page 363,  
10 specifies the priority of modes of operation; they are:

- 1) 100Base-TX, full duplex,
- 2) 100Base-T4,
- 3) 100Base-TX,
- 15 4) 10Base-T, full duplex, and
- 5) 10Base-T.

Two nodes connected via a point-to-point link may negotiate the best common mode of operation between them, according to the above priority. In the prior art, the  
20 nodes are not able to <sup>so</sup> negotiate when a network probe is inserted therebetween. The finite state machine of the present invention, operating at the probe, makes Auto-Negotiation between the nodes possible notwithstanding the physical presence of the probe between the nodes. The finite state machine operating in the probe effectuates the transfer of information of availability of multiple modes of operation between the nodes  
25 in the point-to-point link.



Figure 6 illustrates a block diagram of a probe as may be utilized by an embodiment of the present invention. The probe 100 comprises a plurality of input/output ports, one input/output port 101 for each point-to-point link supported by the probe 100. Each input/output port 101, in turn, is divided into as many channels as there are possible common transmission speeds between the nodes. Each channel comprises a pair of repeaters for transmitting data at a data transmission speed that matches the data transmission speed of the network devices attached to the channel. Two repeaters are necessary to avoid collisions when operating in full duplex mode, i.e., when transmitting and receiving data from the attached nodes at the same time. Thus, for example, probe 100 comprises a channel defined by 10 million bits per second (10 Mb/s) repeaters 110 and 112 for receiving and transmitting data in accordance with the IEEE 802.3 CSMA/CD protocol/Ethernet protocol, as well as a channel defined by 100 Mb/s repeaters 111 and 113 for handling IEEE 802.3u CSMA/CD protocol/Fast Ethernet traffic between the nodes. It is appreciated that other channels/repeaters may be included in a channel depending on the need to support yet other protocols operating at different speeds.

Each repeater transfers data traffic between two physical layer (PHY) interfaces. For example, 10 Mb/s repeater 110 may receive and transmit data between PHYs 120 and 121 in half duplex (Figure 7), or full duplex (Figure 8). If operating in half duplex mode, each channel has one repeater operational between two PHYs, and one receive data channel *and* one transmit data channel operational per PHY. If operating in full duplex mode, each channel has *two* repeaters operational, each between two PHYs, and one receive *or* one transmit data channel operational per PHY.

With reference to Figs. 5 - 10, an embodiment of the present invention is now described. Initially, the probe 100, in which an instance of a finite state machine 148 of



the present invention is implemented, is powered off, and electromechanical relays 160 and 161 are in bypass mode, that is, the relays are closed so that data transmitted by stations 130 and 132 bypass the probe altogether. Upon power up, the PHYs are configured in accordance with normal startup procedures and become operational. The probe opens the relays so that it is in pass through mode, i.e., the relays are opened so that data transmitted by stations 130 and 132 pass through the probe for monitoring. Each PHY to which a station is attached performs Auto Negotiation with the station in accordance with IEEE 802.3u clause 28. For example, PHY 120 performs Auto-Negotiation with station 130, and PHY 121 performs Auto-Negotiation with station 132. After Auto Negotiation is complete, the finite state machine 148 has obtained the mode of operation from each PHY and compares the modes of operation for each station negotiated by the respective PHY to which the station is coupled. The finite state machine then enables the appropriate repeaters and paths between the PHYs to monitor the traffic as it passes through the probe, based on the modes of operation negotiated between the PHYs and the stations to which they are attached.

Probe 100 comprises multiple pass through ports. A single pass through port 101 is illustrated in Figure 6. One pass through port is capable of interconnecting two nodes to form a point to point link that may be monitored by the probe when powered up, or which bypasses the probe if the probe is powered off. A station 130 is initially coupled to PHY 120 via a receive data line 170 and transmit data line 171a. Transmit data line 171a may be switched for transmit data line 172b by switch 150 to transmit data from PHY 123 via line 171b depending on whether station 130 is capable of transmitting and receiving data in full duplex, as determined by Auto-Negotiation. If station 130 operates in full duplex, switch 150 is controlled by finite state machine (FSM) 148 to couple the transmit line 171b from PHY 123 to station 130. If station 130 operates in half duplex,



FSM 148 controls switch 150 to couple the transmit line 171a from PHY 120 to station 130 for half duplex operation.

A station 132, at the other end of the point to point link from station 130, is initially coupled to PHY 121 via transmit data line 173 and receive data line 172a. Receive data line 172a may be switched by switch 152 in favor of receive data line 172b between PHY 122 and station 132 depending on whether station 132 is capable of transmitting and receiving data in full duplex. If station 132 operates in full duplex, switch 152 is controlled by finite state machine (FSM) 148 to couple PHY 122 to station 132 via receive data line 172b. If station 132 operates in half duplex, FSM 148 controls switch 152 to couple PHY 121 to station 132 via receive data line 172a. Thus, switches 150 and 152 control the selection of PHYs with which stations 130 and 132 transmit and receive data depending on the duplex mode of operation of the stations. More particularly, FSM 148 controls PHYs so that full duplex or half duplex communication is established between nodes 130 and 132. If both nodes can support full duplex communication, then full duplex communication is established. If one or both nodes can support only half duplex communication, then half duplex communication is established, as that is the best common mode of operation between the two nodes.

Switches 140-145 are controlled by FSM 148 to utilize certain of repeaters 110-113 depending on the mode of operation of nodes 130 and 132. If, for example, half-duplex, 10 Mb/s transmission speed is the highest common mode of operation between nodes 130 and 132, FSM 148 controls, in addition to enabling appropriate ones of PHYs 120-123 as described above, selecting appropriate ones of repeaters 110-113. In particular, repeater 110 is enabled by selecting receive data line 190a at switch 140, and transmit data line 191a at switch 141, receive data line 192a at switch 142, and transmit data line 193a at switch 143. All other lines at all other switches 140-145 are disabled.



Figure 7 illustrates the circuitry depicted in Figure 6 that is utilized by probe pass through port 101 for exchanging data between nodes 130 and 132 at a particular data transmission speed (10 Mb/s) and in half duplex. For purposes of illustration, the switches that effect the circuitry used are not included in the drawing.

5

If, on the other hand, full duplex, 10 Mb/s transmission speed is the highest common mode of communication between nodes 130 and 132, FSM 148 enables the appropriate paths between PHYs 120-123 and repeaters 110 *and* 112. Both 10 Mb/s repeaters are used to avoid collisions that would otherwise occur if both stations

10 concurrently transmitted data to the same repeater. In particular, receive link 190a between PHY 120 and repeater 110 is enabled by switch 140, transmit data link 193a is enabled between PHY 121 and repeater 110 by switch 142, receive data link 194a is enabled between PHY 122 and repeater 112 by switch 144, and transmit data link 195a is enabled between repeater 112 and PHY 123 by switch 145. The enabled circuitry is  
15 illustrated in Figure 8. Again, for clarity, the FSM and switches controlled thereby are not shown.

Figure 9 illustrates the circuitry enabled by FSM 148 to provide 100 Mb/s half duplex communication between nodes 130 and 132. Switch 140 closes receive data line  
20 190b, while transmit data link 191b is enabled by switch 141, thereby providing for half duplex communication between PHY 120 and repeater 111, in support of station 130. Likewise, FSM 148 controls switches 142 and 143 to links 192b and 193b respectively, to provide half duplex communication between PHY 121 and repeater 111 in support of station 132. All other switches are open, thereby disabling communication over PHYs  
25 122 and 123 and repeaters 112 and 113.



100 Mb/s, full duplex communication between nodes 130 and 132 is depicted in Figures 6 and 10, in which FSM 148 controls switches 140 and 143 to enable receive and transmit data lines 190b and 193b coupled to respective PHYs 120 and 121 for communication with 100 Mb/s repeater 111. FSM 148 further controls switches 144 and 145 to enable receive and transmit data lines 194b and 195b respectively coupled to PHYs 122 and 123 for communication with 100 Mb/s repeater 113. Both repeaters are utilized to avoid collisions that would otherwise occur if stations 130 and 132 concurrently transmitted data to the other station.

As stated above, the relays for each channel are in bypass mode at power down of the probe. Upon power up, FSM 148 is in idle state 505. Configuration of the PHYs, including Auto-Negotiation, occurs next in accordance with the appropriate IEEE standard 802.3 procedures, and the relays are opened so that the probe is in pass through mode. In pass through mode, data packets transmitted between nodes connected to the probe are monitored by the probe, rather than bypassing the probe. At this point, the probe configuration is as illustrated in Figure 7, in which the stations communicate with each other according to the lowest priority mode of operation, e.g., 10 Mb/s data transmission speed, half duplex.

Once Auto-Negotiation is completed by each PHY to which a station is connected, and the probe is in pass through mode, the finite state machine 148 compares the modes of operation negotiated by each PHY with its respective attached station. If only a single station in the point-to-point link, e.g., station 130 or station 132, but not both, is connected to a PHY at 10 Mb/s transmission speed, then FSM 148 transitions to state 510. If both stations in the point-to-point link are connected to a PHY at 10 Mb/s, then FSM transits to state 525, and the probe configuration remains as illustrated in Figure 7.



